

AMENDMENTS TO THE CLAIMS:

This listing of claims will replace all prior versions, and listings, of claims in the application:

Listing of Claims:

1. (Original) A method for managing caches in a system with multiple caches that may contain different copies of a data item, comprising the steps of:
modifying the data item in a first node of said multiple caches to create a modified data item;
sending the modified data item from said first node to a second node of said multiple caches without durably storing the modified data item from said first node to persistent storage;
after said modified data item has been sent from said first node to said second node, said first node sending a request to a master of said data item for writing said data item to persistent storage; and
in response to said request, said master coordinating with said multiple caches to cause said data item to be written to persistent storage.
2. (Original) The method of Claim 1 wherein:
the method includes the step of maintaining, within an ordered series of bins, entries for past-image versions of data items;
each bin in said ordered series corresponds to a time range;
a particular bin corresponds to the time range that covers the time at which the data item is modified in said first node; and
the step of sending a request is performed by sending a request for writing a particular bin of said ordered series of bins to persistent storage.

3. (Original) The method of Claim 2 wherein the step of said master coordinating with said multiple caches to cause said data item to be written to persistent storage includes said master causing said multiple caches to write data items to persistent storage to cover all past image versions of data items that were modified during the time range of said particular bin.
4. (Original) The method of Claim 3 further comprising the step of emptying said particular bin after said multiple caches write data items to persistent storage to cover all past image versions of data items that were modified during the time range of said particular bin.
5. (Original) The method of Claim 4 wherein the step of emptying said particular bin includes the steps of:
discarding entries within said particular bin that are associated with past images that have last-dirtied times within the time range of said particular bin; and
moving to one or more other bins the entries within said particular bin that are associated with past images that have last-dirtied times later than the time range of said particular bin.
6. (Original) The method of Claim 1 wherein the step of sending a request to a master is performed by sending the request to a global lock manager.
7. (Original) The method of Claim 1 wherein the step of sending a request to a master is performed by sending the request to a lock manager that is one of a plurality of lock managers within a distributed lock management system.
8. (Original) The method of Claim 1 further comprising the step of sending from the master, to interested nodes, write-notification messages indicating that said data item

has been written to persistent storage, in response to said data item being written to persistent storage.

9. (Original) The method of Claim 8 wherein the step of sending write-notification messages includes the master sending to at least one interested node a single message that notifies said at least one interested node that a plurality of data items have been written to persistent storage.
10. (Original) The method of Claim 1 wherein the step of said first node sending a request to a master of said data item for writing said data item to persistent storage includes the first node sending to said master a single message that requests writing a plurality of data items to persistent storage, wherein said plurality of data items includes said data item.
11. (Original) The method of Claim 10 wherein the step of sending a single message includes sending a message that identifies a bin to request that all data items that belong to the bin be written to persistent storage.
12. (Original) The method of Claim 11 wherein the bin is associated with a range of time and includes data items that were first dirtied by the first node during said range of time and that were subsequently transferred to other nodes without first being written to persistent storage.
13. (Original) The method of Claim 8 wherein the step of sending from the master to interested nodes write-notification messages includes the steps of:
immediately sending write-notification messages to a first set of interested nodes,
where said first set of interested nodes includes the interested nodes that have requested said data item to be written to persistent storage; and

delaying the sending of write-notification messages to a second set of nodes, where
said second set of nodes includes interested nodes that do not belong to said
first set of interested nodes.

14. (Original) The method of Claim 8 wherein the step of sending from the master to interested nodes includes delaying the sending of write-notification messages to at least one interested node.
15. (Original) The method of Claim 14 wherein a write-notification message is sent to the at least one interested node in response to a lock request made by said at least one interested node.
16. (Original) The method of Claim 14 wherein a write-notification message is sent to the at least one interested node in response to the at least one interested node requesting that said data item be written to persistent storage.
17. (Original) The method of Claim 14 wherein a write-notification message is sent to the at least one interested node within a ping request that the master sends to the at least one interested node for the at least one interested node to transfer another data item to another node.
18. (Original) The method of Claim 1 wherein the step of coordinating includes the steps of:
determining whether a version of said data item, that is at least as recent as said modified version, has already been written to persistent storage; and
if a version of said data item that is at least as recent as said modified version has already been written to persistent storage, then sending a write-notification message from said master to notify said first node that a version of said data

item that is at least as recent as said modified version has already been written to persistent storage.

19. (Original) The method of Claim 18 wherein the step of coordinating includes, if a version of said data item that is at least as recent as said modified version has not already been written to persistent storage, then sending a write-perform message from said master to grant permission for said modified version to be written to persistent storage.
20. (Original) The method of Claim 1 wherein the step of coordinating includes the steps of:
selecting a particular node of said multiple caches that has a particular version of said data item, wherein said particular version is at least as recent as the modified data item in said first node; and
causing said particular version of said data item to be written from said particular node to persistent storage.
21. (Original) The method of Claim 20 wherein the step of selecting a particular node includes selecting the node, of said multiple caches, that has a most recently modified version of said data item.
22. (Original) The method of Claim 20 further comprising the step of the master informing the first node that said data item has been written to persistent storage in response to the master receiving confirmation that said particular version of said data item has been written to persistent storage.
23. (Original) The method of Claim 20 further comprising the step of the master informing a set of caches that said data item has been written to persistent storage in

response to the master receiving confirmation that said particular version of said data item has been written to persistent storage, wherein said set of caches includes caches, other than said particular node, that contain modified versions of said data item that are not more recent than said particular version.

24-48. (Cancelled).

49. (Original) A computer-readable medium carrying instructions for managing caches in a system with multiple caches that may contain different copies of a data item, the instructions comprising instructions for performing the steps of:

modifying the data item in a first node of said multiple caches to create a modified data item;

sending the modified data item from said first node to a second node of said multiple caches without durably storing the modified data item from said first node to persistent storage;

after said modified data item has been sent from said first node to said second node, said first node sending a request to a master of said data item for writing said data item to persistent storage; and

in response to said request, said master coordinating with said multiple caches to cause said data item to be written to persistent storage.

50. (Original) The computer-readable medium of Claim 49 wherein:

the computer-readable medium includes instructions for performing the step of maintaining, within an ordered series of bins, entries for past-image versions of data items;

each bin in said ordered series corresponds to a time range;

a particular bin corresponds to the time range that covers the time at which the data item is modified in said first node; and

the step of sending a request is performed by sending a request for writing a particular bin of said ordered series of bins to persistent storage.

51. (Original) The computer-readable medium of Claim 50 wherein the step of said master coordinating with said multiple caches to cause said data item to be written to persistent storage includes said master causing said multiple caches to write data items to persistent storage to cover all past image versions of data items that were modified during the time range of said particular bin.
52. (Original) The computer-readable medium of Claim 51 further comprising instructions for performing the step of emptying said particular bin after said multiple caches write data items to persistent storage to cover all past image versions of data items that were modified during the time range of said particular bin.
53. (Original) The computer-readable medium of Claim 52 wherein the step of emptying said particular bin includes the steps of:
discarding entries within said particular bin that are associated with past images that have last-dirtied times within the time range of said particular bin; and
moving to one or more other bins the entries within said particular bin that are associated with past images that have last-dirtied times later than the time range of said particular bin.
54. (Original) The computer-readable medium of Claim 49 wherein the step of sending a request to a master is performed by sending the request to a global lock manager.
55. (Original) The computer-readable medium of Claim 49 wherein the step of sending a request to a master is performed by sending the request to a lock manager that is one of a plurality of lock managers within a distributed lock management system.

56. (Original) The computer-readable medium of Claim 49 further comprising instructions for performing the step of sending from the master, to interested nodes, write-notification messages indicating that said data item has been written to persistent storage, in response to said data item being written to persistent storage.
57. (Original) The computer-readable medium of Claim 56 wherein the step of sending write-notification messages includes the master sending to at least one interested node a single message that notifies said at least one interested node that a plurality of data items have been written to persistent storage.
58. (Original) The computer-readable medium of Claim 49 wherein the step of said first node sending a request to a master of said data item for writing said data item to persistent storage includes the first node sending to said master a single message that requests writing a plurality of data items to persistent storage, wherein said plurality of data items includes said data item.
59. (Original) The computer-readable medium of Claim 58 wherein the step of sending a single message includes sending a message that identifies a bin to request that all data items that belong to the bin be written to persistent storage.
60. (Original) The computer-readable medium of Claim 59 wherein the bin is associated with a range of time and includes data items that were first dirtied by the first node during said range of time and that were subsequently transferred to other nodes without first being written to persistent storage.
61. (Original) The computer-readable medium of Claim 56 wherein the step of sending from the master to interested nodes write-notification messages includes the steps of:

immediately sending write-notification messages to a first set of interested nodes,
where said first set of interested nodes includes the interested nodes that have
requested said data item to be written to persistent storage; and
delaying the sending of write-notification messages to a second set of nodes, where
said second set of nodes includes interested nodes that do not belong to said
first set of interested nodes.

62. (Original) The computer-readable medium of Claim 56 wherein the step of sending from the master to interested nodes includes delaying the sending of write-notification messages to at least one interested node.
63. (Original) The computer-readable medium of Claim 62 wherein a write-notification message is sent to the at least one interested node in response to a lock request made by said at least one interested node.
64. (Original) The computer-readable medium of Claim 62 wherein a write-notification message is sent to the at least one interested node in response to the at least one interested node requesting that said data item be written to persistent storage.
65. (Original) The computer-readable medium of Claim 62 wherein a write-notification message is sent to the at least one interested node within a ping request that the master sends to the at least one interested node for the at least one interested node to transfer another data item to another node.
66. (Original) The computer-readable medium of Claim 49 wherein the step of coordinating includes the steps of:
determining whether a version of said data item, that is at least as recent as said
modified version, has already been written to persistent storage; and

if a version of said data item that is at least as recent as said modified version has already been written to persistent storage, then sending a write-notification message from said master to notify said first node that a version of said data item that is at least as recent as said modified version has already been written to persistent storage.

67. (Original) The computer-readable medium of Claim 66 wherein the step of coordinating includes, if a version of said data item that is at least as recent as said modified version has not already been written to persistent storage, then sending a write-perform message from said master to grant permission for said modified version to be written to persistent storage.
68. (Original) The computer-readable medium of Claim 49 wherein the step of coordinating includes the steps of:
selecting a particular node of said multiple caches that has a particular version of said data item, wherein said particular version is at least as recent as the modified data item in said first node; and
causing said particular version of said data item to be written from said particular node to persistent storage.
69. (Original) The computer-readable medium of Claim 68 wherein the step of selecting a particular node includes selecting the node, of said multiple caches, that has a most recently modified version of said data item.
70. (Original) The computer-readable medium of Claim 68 further comprising instructions for performing the step of the master informing the first node that said data item has been written to persistent storage in response to the master receiving

confirmation that said particular version of said data item has been written to persistent storage.

71. (Original) The computer-readable medium of Claim 68 further comprising instructions for performing the step of the master informing a set of caches that said data item has been written to persistent storage in response to the master receiving confirmation that said particular version of said data item has been written to persistent storage, wherein said set of caches includes caches, other than said particular node, that contain modified versions of said data item that are not more recent than said particular version.

72-96. (Cancelled).